

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

TITLE OF THE INVENTION

**CLOSING A COMMUNICATIONS STREAM BETWEEN TERMINALS OF A
COMMUNICATIONS SYSTEM**

INVENTORS

DOUGLAS C. DAHLBY

MITCHELL D. TROTT

CHRISTOPHER RICHARD UHLIK

Prepared by

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026
(303) 740-1980

EXPRESS MAIL CERTIFICATE OF MAILING

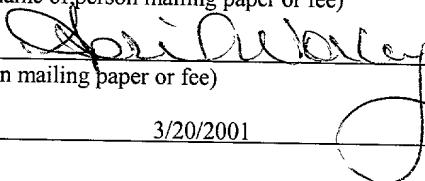
"Express Mail" mailing label number: EL752329388US

Date of Deposit: March 20, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20231

April M. Worley

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

3/20/2001

(Date signed)

CLOSING A COMMUNICATIONS STREAM BETWEEN TERMINALS OF A COMMUNICATIONS SYSTEM

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates generally to the field of digital signal communications and channel allocation. More particularly, the invention relates to closing a communications stream between two communicating terminals using a very short messaging sequence with minimum latency.

Description of the Related Art

[0002] In data communications systems that share communications channels, a procedure must be implemented to allow different users to take turns using the channels. In many systems, users are allocated a channel for a limited period of time. Then the channel is de-allocated and allocated to another user. This is known as a circuit-switched system. In circuit-switched systems extensive setup must be done so that subsequent packets do not need to carry the full routing information. The packets instead will always travel the same circuit until that circuit is switched. In other systems, channels are never allocated but are always shared. This kind of system requires no overhead signaling for channel allocation and deallocation but must have some way of directing packets and of resolving conflicts. Such systems are typically called packet-switched because each packet contains the full routing information required to direct the packet to its intended destination. Packet switching requires no overhead in the setup but adds significant overhead to each packet.

[0003] In order to maximize the use of shared channels, the system should be able to efficiently use channels for very short periods. The system may be inefficient if the time needed to allocate and deallocate the channel is large in comparison to the time that the channel is in use.

[0004] The present invention relates to two major aspects of channel deallocation. The first is the decision of when to deallocate a channel; the second is how to perform the deallocation as efficiently as possible. The channel deallocation decision involves a tradeoff between minimizing system usage efficiency and minimizing data latency for individual users. In some cases, the deallocation decision can benefit both resource usage efficiency and latency, but in most cases the decision will improve one and degrade the other.

[0005] This trade-off has been addressed in previous systems in several ways. In one approach, channels are deallocated with a multiple step handshake. The handshake includes an explicit acknowledgement in order to deallocate the channel. This requires that the initiator of the deallocation keep the channel open until the peer has responded. It also requires that the peer keep the channel open long enough to send the response. This reduces the amount of time that the channel can be available for others. In other systems, still more stages are required in that there are multiple interactions. This may happen, for example, if multiple protocol layers are involved in the allocation and deallocation of each channel, so each protocol layer has to be terminated separately.

[0006] In other systems, a channel is terminated immediately when there is no data, in spite of the finite channel allocation and deallocation cost. This can waste system resources and increase user latency. On the other hand to keep a channel open for a long period when there is no data also wastes system resources. An alternative approach is not

to allocate channels at all; instead all users share channels by putting full routing information into every packet, a packet-switched system. In a packet-switched system, there is no allocation or deallocation, but efficiency is reduced because there is extra overhead for every packet.

[0007] The present invention can minimize the difficulty of closing communications channels. This allows channels to be used for very short durations without dedicating a large amount of the communications resources to overhead. As a result, the communications system can be packet oriented rather than circuit oriented without being packet-switched. As a result, the communications system can efficiently create circuits which only deliver a few packets. The present invention also can optimize the amount of time during which a channel is left open even when there is no data to transmit.

[0008] Among the difficulties to be minimized in closing communications that can be minimized are latency and interference between users. Latency between the time a channel is no longer needed and the time it is closed can be reduced by optimizing the channel deallocation messages. Latency between the time a channel is no longer needed and the time that a decision to deallocate is reached can be minimized by providing other channel characteristics and parameters along with the channel request. It can also be minimized using channel characteristics and smart decision parameters that have already been established by channel and data history. Limits on system capacity due to contention between users can be reduced by limiting the number of messages required to deallocate a channel and by reducing the number of allocation and deallocation cycles.

BRIEF SUMMARY OF THE INVENTION

[0009] A method and apparatus are provided that quickly and efficiently closes a stream between two communicating terminals that use repeating time slots in a time division channel. In one embodiment, the invention includes sending a closing message in a first slot from the first terminal to the second terminal to request that the communications stream be closed, listening to the communications stream at the first terminal to determine whether any messages are sent from the second terminal to the first terminal in a slot after the first slot, and closing the stream, if no further messages are received from the second terminal.

[0010] In another embodiment, the invention includes checking a transmit buffer for data to transmit in the stream, waiting for a determined amount of time, rechecking the transmit buffer, and determining to close the communications stream, if there is no data in the transmit buffer to transmit after the checking and the rechecking.

[0011] Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0012] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

[0013] **Figures 1A and 1B** are a diagram showing an exchange of messages between two terminals in accordance with the present invention;

[0014] **Figure 2** is a flow diagram showing a decision process for closing a stream in accordance with the present invention;

[0015] **Figure 3** is a block diagram illustrating an exemplary architecture of a wireless communication system appropriate for use with one embodiment of the present invention; and

[0016] **Figure 4** is a block diagram illustrating an alternate exemplary architecture of a wireless communications system appropriate for use with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0017] Introduction

[0018] The present invention allows for the communication of data, which may represent numerical, textual, graphical, pictorial, audio, video, voice or any other type of information between terminals in a communications system. Generally, the terminals will be referred to herein as user terminals and base stations and communications will be between a user terminal and a base station. However, the terminals can be in any of a variety of other configurations and communications can be between peers, such as two user terminals or two base stations and communications can be broadcast or multicast as is desired for the particular system under consideration.

[0019] The data is transmitted as packets in duplexed circuit-switched channels called streams. Each new channel constituting a new stream. The stream, in one embodiment, is a slot assignment of a frequency-hopping time division duplex TDMA (time division multiple access) radio communications system, although other types of channels are also appropriate for application of the present invention depending upon the circumstances. In this embodiment, the stream is made up of an uplink slot to send data to the base station and a downlink slot to receive data from the base station. The slots are a portion of a repeating 5msec frame on a channel that hops frequency every frame.

[0020] In one embodiment, a stream is opened when a user terminal transmits a request to the base station that will be called RA-rts (Random Access - request to send). The RA-rts is sent on a previously designated channel for RA (Random Access) messages. In one embodiment the RA channel is a set of channels, which may be the

same set of channels used to exchange traffic bursts. The base station, upon receiving the message, determines whether a stream can be opened for the requesting user terminal and, if so, sends an AA - cts (Access Acknowledgment - clear to send) message back to the user terminal that identifies the channel for the stream, both the uplink portion on which the user will send and the downlink portion on which the user terminal will receive. This exchange opens the stream and allows data to be exchanged. In one embodiment, in-band signaling is provided so that further signaling, requests and control messages can also be sent in the stream thereby freeing capacity in the random access channels. Because of a prior registration process to be described in more detail below, no further messages are necessary in order to open a stream. As a result, the stream can be opened within the course of one uplink random access timeslot and the following downlink timeslot within the same frame. Assigning streams quickly allows streams to be closed more frequently without significantly affecting communications speed between the user terminal and the base station and without causing the channel to be used for stream assignment when it could be used for stream data traffic. This allows the total capacity of shared system resources to be increased.

[0021] The user terminal will typically open a stream when it has data in a buffer to be sent to the base station. As mentioned above, the data may be of any kind including a request for data from or via the base station. If, for example, the user terminal is being used as a web browser, the data can be simply an HTML (Hyper-Text Markup Language) request to retrieve a web page with a particular URL (Uniform Resource Locator). Such a message when received by the base station might be transmitted to an ISP (Internet Service Provider) to retrieve the page or the base station may be an ISP.

[0022] The user terminal may also request that a stream be opened in response to a request from a base station. In one embodiment, the base station can send an RA-rts to the user terminal, which can then respond with AA-rts. In another embodiment, the base station can send a page message to the user terminal asking the user terminal to request that a stream be opened. If the user terminal is able to, then in response to the page message, the user terminal sends a RA-rts message to the base station, as discussed before. The advantage of initiating the stream through the user terminal will become clearer as the various messages are described below. Requiring the base station to make all determinations regarding opening streams allows communications system resources to be more efficiently utilized under the central control authority of the base station.

[0023] The user terminal or the base station will typically close a stream when it has no more data in its buffer to be sent or shortly thereafter. The request to close the stream is sent in the data traffic channel of the stream as in-band signaling. If, for example, the user terminal is being used as a web browser and a requested web page has been received, then the base station may request that the stream be closed. If another web page is requested from the base station, then the user terminal can open a stream to send the request and to retrieve the next page. If the pages are requested with very little time in between the requests, then the stream may be kept open.

[0024] In one embodiment, either the user terminal or the base station can send an alert to the other terminal, which can then be responded to with an immediate shutdown. Any remaining data to be sent will be held for the next opened stream and if no stream is opened, then it will eventually be discarded. In another embodiment, either the user terminal or base station can send a shut down request to the other terminal asking the other terminal to close the stream. If the other terminal has no further data to transmit,

then in response to the shut down request message, it will stop transmitting and the stream will be closed. If the other terminal has data left to transmit, then the requesting terminal may send the shut down request several times before the other terminal's transmit data buffer is exhausted and the stream is closed.

[0025] In one embodiment, the present invention is implemented in a Spatial Division Multiple Access (SDMA) radio data communications system as shown and described further with respect to Figures 3 and 4. In such an SDMA system, each terminal is associated with a set of spatial parameters that relate to the radio communications channel between, for example, the base station and a user terminal. The spatial parameters comprise a spatial signature for each terminal. Using the spatial signature and arrayed antennas, data can be sent to several different user terminals at the same time using the same radio frequency and time slot. Conversely, data can be received from several different user terminals at the same time on the same radio frequency and time slot. The spatial signatures allow the otherwise colliding signals to be resolved. The spatial signatures can include such things as the spatial location of the transmitters, (the directions-of-arrival (DOAs), times-of-arrival (TOAs) and the distance from the base station. Estimates of parameters such as signal power levels, DOAs, and TOAs can be determined using known training sequences placed in digital data streams for the purpose of channel equalization in conjunction with sensor (antenna) array information. This information is then used to calculate appropriate weights for spatial demultiplexers, multiplexers, and combiners. Extended Kalman filters or other types of linear filters, well-known in the art, can be used to exploit the properties of the training sequences in determining spatial parameters. Further details regarding the use of SDMA

systems are described, for example, in U.S. Patents Nos. 5,828,658, issued Oct. 27, 1998 to Ottersten et al. and 5,642,353, issued June 24, 1997 to Roy, III et al.

[0026] Registration

[0027] A user terminal forms a relationship with a base station called a registration or session. This registration begins by listening to a BCH (Broadcast Channel) and ends with a handover, timeout, or disconnection. The first step of registration is accomplished by a user terminal by sending a CR (Configuration Request) burst and receiving a CM (Configuration Message) burst. The CM contains basic configuration parameters such as hopping sequence calculation parameters. Using the information from the CM, the user terminal then opens an unauthenticated registration stream using a RA-req (Random Access registration request) message. This unauthenticated stream carries only in-band signaling data used to complete registration and assignment of a RID (Registration Identifier) and PID (Paging Identifier). Using the RID, assigned at the end of the registration stream, the user terminal can open subsequent streams and it can end the registration. Using the PID, the user terminal can identify the pages that are directed to it from the base station. A user terminal can have more than one registration with the same or several different base stations. Each registration and session has a unique RID.

[0028] During the registration stream, identities and capabilities are exchanged, operating parameters are set, and a RID and PID are assigned. Later, streams may be created and attached to this RID, or the registration state, i.e. RID, PID, and operating parameters, may be replaced by another stream on another base station or another base

station modem on the same base station using the same RID or PID. The specific details of registration are provided here as examples only. Many other registration scenarios are also possible within the scope of the present invention.

[0029] Referring to Figures 1A and 1B, the frame timing is established by the base stations that are in the area and transmitting on a pre-programmed RF carrier. The carrier may be a frequency hopping or spread spectrum carrier. However, it is preferred that the carrier be easy to find and be pre-programmed into the user terminals. The base stations, or base station if there is only one, employ GPS or some other precise common timing reference to establish the frame timing 300. GPS timing offers the advantage that it is accurately synchronized and inexpensively available to all base stations. This allows the BCH to be shared by all the base stations with only a minimal guard time in the BCH between base stations. The base stations then build the BCH frame 302 described above and broadcast in their respective assigned slots 304.

[0030] When a user terminal turns on, it scans the BCH carrier, a well-known, optionally pre-programmed, RF carrier 306 to find basic frame timing 308 and synchronization 310. The user terminal scans this carrier for BCH bursts, building an RSSI (Received Signal Strength Indicator) map 312. From this BCH RSSI map and other factors, the user terminal selects the strongest or the best base station 314. It also uses the BCH to precisely adjust its oscillator frequency and adjust its frame timing reference 308, 310. This is done using the synchronization and timing sequences in the BCH burst. Then, using its UTID (User Terminal Identification) it builds 316 and sends 320 a CR (Configuration Request), timed relative to the BCH burst for that strongest or best base station. In one embodiment, the CR is scrambled using the BSCL that was received in the BCH from the selected base station 318.

[0031] If the intended base station successfully receives the CR and has available capacity, it unscrambles the CR 322, and determines the spatial signature of the user terminal 324. The user terminal receives a CM (Configuration Message) burst in reply 326. The CM contains sufficient information for the user terminal to learn its distance and RF path-loss to the base station, correct its timing advance, adjust its power control, and learn the parameters of frequency hopping (e.g. frame numbering and BSCC (Base Station Color Code) 328. Several base stations may be probed with a CR to find the closest or the best base station. Based on this information from the CM, the user terminal, when it has data to transmit, can start a session, beginning with a RA-rreq (Random Access - registration request) 330. If resources are available, the base station sends an AA - reg-ack (Access Assignment - registration-acknowledgment) 332 to the user terminal assigning a traffic channel for the registration process. The base station and user terminal exchange various access control parameters including encryption keys on this established stream. Finally a RID and PID are assigned. Using this RID, the user terminal can establish secure streams (using e.g. RA-rts/AA-cts as discussed above) in which it transmits and receives data packets 334 until the registration stream is closed 344.

[0032] The traffic channel includes a data acknowledgement DA or a data invalid DI response to each transmission. The DA and DI messages are transmitted in the transmission from the recipient in the next timeslot. In a time division duplex frame, the base station and the user terminal alternate slots 336, 338, 340, 342 or group of slots, as shown in Figure 1A. Accordingly, if any slots are not received properly, the data can quickly be retransmitted. This reduces the size of the data buffers at the respective base station and user terminal modems. As shown in Table 1, uplink slots always precede

downlink slots and there is a guard time between the two in order to allow for any synchronization errors or unanticipated propagation delays. In one embodiment, each side transmits data packets in three slots, each slot including ramp-up and ramp-down periods as well as synchronization bits as is well-known in the art.

| | | | | | | | | | | | |
|--------------|------------|----------------|------------|---|--------------|---|--|---|---|---|-----|
| 1 | 2 | 3 | | 1 | 2 | 3 | | 1 | 2 | 3 | ... |
| Uplink Slots | Guard Time | Downlink Slots | Guard Time | | Uplink Slots | | | | | | |

Table 1

[0033] Periodically, the user terminal scans the BCH to update its RSSI and BSCC map. When it detects a better base station, it may send a CR to this new base station and possibly handover its network session, i.e. its registration. If successful stream initiation fails too many times, the user terminal enters a timeout state. From timeout, it may try to regain a RID via RA-req, refresh its timing advance using a CR, find a new base station to which it might handover by scanning the BCH, or even begin from scratch to re-acquire basic frame timing. If this re-establishment is successful, the user terminal may be able to continue its network session by completing a network session handover to the new base station.

[0034] Configuration Request CR

[0035] Table 2 below summarizes the content of the example CR burst. The 82 information symbols are constructed from the configuration request message using modulation and coding.

| Duration | Contents |
|---------------|---|
| 10 μ sec | ramp-up |
| 260 μ sec | training symbols a_1, a_2, \dots, a_{130} |
| 164 μ sec | information symbols h_1, h_2, \dots, h_{82} |
| 10 μ sec | ramp-down |
| 86 μ sec | extra guard time |
| 15 μ sec | inter-burst guard time |

Table 2

[0036] CR spatial training is the same for all base stations and the base station does not necessarily know the location of the user terminal before receiving the CR. CRs are transmitted by user terminals at a fixed offset from BCH transmissions. CR and CM are scrambled by a function of BSCC ensuring that even if there is some interference from CRs sent to nearby base stations, the demodulation capture effect of the BSCC works out any collisions. In one embodiment, the scrambling is performed by taking the encoded bit sequence and exclusive OR'ing it with the output of a linear feedback shift register. Finally the smart antenna spatial resolution ability of the base station is applied to resolve any remaining ambiguities in received CRs.

[0037] A configuration request message is mapped onto a configuration request burst by the physical layer. A configuration message is mapped onto a standard downlink burst by the physical layer. The information symbols of the CR burst are mapped out as shown in Table 3, below. Any of the items listed below can be deleted and transmitted later during the registration cycle or not at all based on the needs of the system.

| Configuration Request Message | |
|-------------------------------|-----------|
| Field | # of Bits |
| identity | 8 |
| utClass | 4 |
| txPwr | 5 |
| Total | 17 |

Table 3

[0038] The meanings of the symbol sets are as follows:

identity: a set of unique random bits for each user terminal that differentiate simultaneous messages from multiple user terminals. Because of the randomness and large number of bits, it is unlikely that two user terminals will select the same identity code at the same time.

utClass: identifies user terminal capabilities (highest modulation class, frequency hopping capabilities, etc.). This sequence identifies the type of user terminal that sent the CR. A palmtop digital assistant might have different capabilities than a desktop computer with a fixed dedicated antenna. With utClass, the different capabilities can be distinguished.

txPwr: the power used by the user terminal to transmit the Configuration Request burst. For example, user terminal power = $(2 \cdot \text{txPwr} - 30)$ dBm.

[0039] CR is sent on the control carrier, as an example, exactly 2265 μ sec after receipt of a downlink BCH burst. In this way, an otherwise uninitialized user terminal can send CR without any knowledge of the frequency hopping sequence parameters. The CR burst is shorter than a standard uplink time slot to allow for unknown time-of-flight from

the user terminal to the base station and typically arrives late in the uplink time slot receive window.

[0040] Configuration Message CM

[0041] Table 4 below summarizes the content of an example Configuration Message burst. The 494 information symbols are constructed from the configuration message using modulation and coding.

| Duration | Contents |
|---------------|--|
| 10 μ sec | ramp-up |
| 68 μ sec | training symbols a_1, a_2, \dots, a_{34} |
| 920 μ sec | information symbols h_1, h_2, \dots, h_{460} |
| 68 μ sec | training symbols a_1, a_2, \dots, a_{34} |
| 10 μ sec | ramp-down |
| 15 μ sec | inter-burst guard time |

Table 4

[0042] The configuration message CM burst is sent on the BCH carrier, after sending a downlink BCH burst, whenever CR was received on the corresponding uplink time slot. CM is also sent in a spatially directed signal based on the analysis of the spatial signature of the received CR burst, for example parameters, such as DOA (direction of arrival) and TOA (time of arrival) of the uplink CR. Since CM is sent on the BCH carrier, a fixed time offset from BCH, an otherwise uninitialized user terminal can receive CM without any knowledge of the frequency hopping sequence parameters. CM, in response to CR, includes, among other things; the AFN (Absolute Frame

Number), a larger timing-advance adjustment dynamic range, coarser power control, and various access control parameters. Table 5, below summarizes the content of the CM burst. Any of the items listed below can be deleted and transmitted later during the registration cycle or not at all based on the needs of the system.

| Configuration Message | |
|-----------------------|-----------|
| Field | # of Bits |
| identity | 8 |
| pwrCtrl | 4 |
| timingAdjust | 7 |
| AFN | 10 |
| carrierMask | 16 |
| racarrierMask | 16 |
| raslotMask | 3 |
| raDec | 3 |
| hopping | 1 |
| Total | 70 |

Table 5

[0043] The meanings of the symbol sets are as follows:

identity: the random identity sent by the user terminal in the CR

pwrCtrl: power offset that the user terminal should apply to future parameter request bursts and random access bursts: offset = $(2 \cdot \text{pwrCtrl} - 16)$ dB.

timingAdjust: timing advance that the user terminal should apply to future random access bursts: timing advance = timingAdjust μ s.

AFN: the 10 least significant bits of the Absolute Frame Number

carrierMask: bitmap of carriers containing traffic channels

racarrierMask: bitmap of carriers containing random access channels (least significant bit is carrier 0)

raslotMask: bitmap of slots containing random access channels (least significant bit is slot 1). Random access channels occur where both racarrierMask and raslotMask are nonzero.

raDec: AFNs available for random access channels.

hopping: if equal to 1, the relationship between physical and logical carriers hops each frame.

[0044] Random Access-request to send burst

[0045] As can be seen from the discussion above, after registration, the user terminal has a RID and a PID and a fair amount of information about the network, including all of the data listed in Table 5. This information includes an assigned random access channel or set of assigned random access channels and an initial transmit power level. This information is used in generating and sending the RA-rts.

[0046] Referring again to Figures 1A and 1B, after the user terminal has been registered with a particular base station, it can open a stream for a data exchange. The opening of a stream can be initiated by either the base station or the user terminal. Typically a stream will be opened if either the base station or the user terminal has data to send to the other. This data is buffered until a preset amount has accumulated in a transmit buffer or until a preset amount of time has elapsed. The preset amount can be any non-zero value. If the base station has accumulated transmit data in its buffer 346 for the user terminal, then it will send a page, described in more detail below, to the user

terminal 347. If the user terminal receives a page or if it has accumulated a sufficient amount of data in its transmit buffer 348, then it will send, for example, an RA-rts message 349. This message, as explained below, is a request that a stream be opened to allow the exchange of data. The base station upon receiving the RA message will analyze its system resource availability 350 and if a suitable channel is available, then it will respond with, for example an AA-cts message 351. This message, as explained below, identifies a channel and assigns it for the stream.

[0047] With the RA/AA exchange, the stream is assigned and all the necessary information for the terminals to communicate has been exchanged. With the next uplink slot, the remote terminal will begin sending its data over the assigned channel 352. If the stream was initiated by a page from the base station, then the remote terminal may not have any data to send in which case it will send idle bits. The idle bits help the base station to maintain its spatial parameters for the user when there is no data received. The base station will use these spatial parameters to send its data packets or idle bits 354. In this way, data and acknowledgments are exchanged in the same way as for the registration stream 352, 354, 356.

[0048] Table 6 below summarizes the content of an example Random Access Message burst. The burst structure is the same as an uplink data burst on a traffic channel TCH. For an uplink data burst, the information symbols carry data or in-and signaling or both.

| Duration | Contents |
|---------------|--|
| 10 μ sec | ramp-up |
| 146 μ sec | training symbols a_1, a_2, \dots, a_{73} |
| 364 μ sec | information symbols h_1, h_2, \dots, h_{182} |
| 10 μ sec | ramp-down |
| 15 μ sec | inter-burst guard time |

Table 6

[0049] The RA burst information symbols, in one embodiment, have the fields as shown below in Table 7.

| Random Access Message | |
|-----------------------|-----------|
| Field | # of Bits |
| RAType | 3 |
| ID | 15 |
| UTTxPwr | 5 |
| Total | 23 |

Table 7

[0050] The meanings of the symbol sets are as follows:

RAType: the type of the RA burst as described in association with Table 8.

ID: the registration identifier, either the RID or, for page responses, the PID. This field can be used by the base station to prioritize stream requests. User terminals with a higher priority can be identified by the RID or PID and be granted a stream in preference to other users. The ID is also used to access the registration account and information of the requesting user terminal.

UTTxPwr: the power used by the user terminal to transmit the burst.

[0051] Any one or more of these fields may be deleted or modified and more fields may be added to suit particular applications.

[0052] The RAType field allows for there to be different types of RA messages sent on the same channel. Table 8 lists possible examples that can be supported with a three bit field. Further or different types of RA messages can be used depending upon the particular nature of the network. More bits can be used to allow for still more different types of messages. As an alternative, the user terminal may send a different RA burst depending on the circumstances as listed in Table 8. All of the RA bursts in Table 8 are sent on the random access channel assigned to the user terminal by the base station. In one embodiment, the RA channels are a set of channels that are also used for traffic.

| Value | Symbol | Meaning |
|-------|------------------|-------------------------------------|
| 000 | RA-rts | stream request |
| 001 | RA-ping | keep-alive poll request |
| 010 | RA-rts-short | short stream request |
| 011 | RA-rts-directed | directed stream request |
| 100 | RA-page response | stream request due to page |
| 101 | RA-rts-UM | stream request, unacknowledged mode |
| 110 | RA-rreq | registration request |

Table 8

[0053] The meanings of the symbol sets are as follows:

RA-rts will be discussed further below and is the mechanism with which the user terminal can open a new communications stream following registration.

RA-ping: can be used to alert the base station to the location, channel characteristics and activity of a user terminal without opening a stream. Pinging the base station can be used to keep a registration alive.

RA-rts short, -directed and -UM: can be used to open special types of streams.

RA-page response: can be sent when the user terminal has no data to transmit but is requesting that a stream be opened in response to a page from the base station. In some systems, it may be preferred that the base station open the stream directly without first paging the user terminal as discussed above.

RA-rreg: can be used to open a new registration or change an existing registration. As mentioned above, a user terminal uses the RA bursts after registration, however, it may be useful in network management for a single user terminal to have two registrations for different individuals, different accounts, different types of communications or other reasons.

[0054] Access Assignment Burst

[0055] The user terminal transmits any random access message such as RA-rts on the uplink side of the random access channel. The base station uses the downlink portion of the random access channel to grant random access requests and to assign resources to the requested data stream using an AA (Access Assignment) message. The AA message can have different formats. One format is shown in Table 9 below.

| Access Assignment Message | |
|---------------------------|-----------|
| Field | # of Bits |
| ID | 15 |
| AAType | 3 |
| modClassUp | 5 |
| modClassDown | 5 |
| frameDec | 3 |
| resource ibChan | 6 |
| pwrCtrl | 4 |
| timingAdjust | 5 |
| tOffset spChan | 3 |
| Total | 49 |

Table 9

[0056] The meanings of the symbol sets are as follows:

ID: the id of the user terminal, either the RID or PID that was transmitted in the RA-rts.

modClassUp: identifies the modulation and coding used for the uplink.

modClassDown: identifies the modulation and coding used for the downlink.

frameDec: defines fractional rate channels.

resource ibChan: indicates the uplink/downlink resource pair that is assigned to the stream.

pwrCtrl: a power adjustment for the UT to apply to subsequent transmissions.

timingAdjust: a timing adjustment for the UT to apply to subsequent transmissions.

tOffset: training sequence offset adjustment for the UT to apply to subsequent transmissions.

AAType: indicates the type of Access Assignment message. Many different possible types are possible. Table 10 below provides one example set of AA types.

| Value | Symbol | Meaning |
|-------|-------------------|--------------------------------------|
| 000 | AA-cts | stream grant |
| 001 | AA-reject | request rejected |
| 010 | AA-ping-ack | keep-alive poll acknowledgment |
| 011 | AA-cts-short | short uplink grant |
| 100 | AA-cancel | cancel previous false page |
| 101 | AA-prev-short-ack | previous short uplink was successful |
| 110 | AA-invalid-ack | received RA is not valid |
| 111 | AA-req-ack | registration grant |

Table 10

[0057] The meanings of the symbol sets are as follows:

AA-cts: (Access Assignment-clear-to-send) begins a stream with the sending user terminal based on the parameters in the AA-cts message. An AA-cts can be sent in response to any of the RA messages and is particularly appropriate for RA-rts, RA-ping, RA-rts-directed, RA-page-response, and RA-rreg. This allows the base station to open a stream even if the user terminal was not aware of the need to open a stream. The next communication will be data in the opened stream. As mentioned above, data will be transferred until the corresponding stream data buffers are emptied. Typically the stream

is then closed. However, the stream may be closed upon the occurrence of many other events as well.

AA-reject: can be used to reject the request and direct the UT to start a timer before sending an RA message. Such a response can relieve congestion on a busy base station. The UT in response may elect to wait or to send an RA-rts to another base station that has better traffic availability.

AA-ping-ack: acknowledges the RA-ping and resets the timers for the registration. The pinging process can be used to prevent a registration from expiring for lack of traffic. Maintaining the registration allows a stream to be opened immediately with e.g. an RA-rts and an AA-cts. If the registration expires then the registration process must be repeated before a data stream can be opened.

AA-cts-short, and AA-prev-short-ack: can be used for special types of streams.

AA-cancel: can be used to respond to an RA-page-response when no page was sent or the paging condition no longer applies.

AA-invalid-id: can be used to notify the UT that it is using a RID or PID that has expired or is not valid with the responding base station. The UT can use the information in the AA to request that a new registration stream be opened by sending e.g. RA-rreq.

AA-reg-ack: is the acknowledgment of RA-rreq that begins a registration stream.

[0058] As mentioned above, the base station can send a page to the UT directing the UT to send an RA-page-response message to the base station. In one embodiment, this reduces the control traffic overhead by using a paging channel. The paging channel can be efficiently utilized by the base station and the page can allow random access channel assignments that increase the channel efficiency of establishing the desired data

stream. The page bursts are transmitted on a paging channel which may be used exclusively for pages or it may be shared with other functions such as a broadcast channel or a control channel. Alternatively a section of a traffic channel can be used for pages. Pages will contain an indication of the base station sending the page and of the user terminal being paged, typically the PID. If the UT is already registered then the page does not need to also include any information about how to respond to the page because this information can be included in the registration data exchange stream. In the embodiment discussed above, the UT will respond to a page in the random access channel by sending a RA-page-response message, however other types of responses are also possible.

[0059] In-Band Signaling

[0060] Within the data stream, any uplink or downlink burst can contain tagged data fields. The tagged fields can have user data and in-band signaling data. The in-band signaling can be used for layer 2 and layer 3 messages, including a request to shut down the stream, a request to open another stream, a request to change the type of the stream, or a request to update registration information. The in-band signaling can also be used for other things such as modulation and transmission channel data reports.

[0061] Stream Closing Handshake

[0062] When both sides have exhausted their transmit buffers the stream will be closed 358. Either the base station or the remote terminal can close the stream. The base

station can close the stream for resource management reasons, even if the transmit buffers have not been exhausted.

[0063] In one embodiment, either the base station or the user terminal can close the stream by sending the following messages:

alert: this message requests an immediate closing of the stream. The message can be replied to with no response. When the requesting terminal receives no data on the next traffic channel slot, it can assume that the alert was received by the other station. Such a protocol minimizes the time and overhead required to close the stream.

no new packets: this message requests that the receiving terminal stop adding packets to its transmission buffers. After the currently buffered packets are transmitted, the stream will be shut down. If, later on, further new packets arrive to be sent, then a new stream can be opened. Such a command can allow e.g. a base station to reallocate channels in order to optimize loading from different users.

shut down request: this message signals a request to close the stream. If the receiving terminal has no data to transmit, then it may close the stream by sending no more bursts. If it still has data to transmit, then it may continue to transmit until its buffers are emptied. The shut-down request may be repeated several times by the requesting terminal before the stream is closed.

[0064] The process for a shut down request is illustrated in more detail in Figures 1A and 1B. In the example of Figures 1A and 1B, the base station has made a determination to close the stream, however, the remote terminal may also make the decision in which case the same process would apply but with the roles of the two terminals reversed. The base station sends a shut-down request to the user terminal 358.

The user terminal receives the request, however, in the example of Figures 1A and 1B, it still has data to send, so instead of closing the stream it sends a data burst with an acknowledgment of the last packet from the base station 360. The base station may continue to send shut-down requests 358 which are answered by data bursts 360 for as long as the user terminal desires to send data. If the base station or the user terminal requires a faster closing of the stream, then it can send an alert or a no new packets request at any time. If the shut-down request is used, then the terminal that receives the request will continue to send data bursts until its transmit buffer is depleted and it has no more data to send.

[0065] The base station continues to send the shut-down requests 362 and after the user terminal's transmit buffer is depleted, it closes the stream on its side and stops transmitting 365. The base station continues to listen to the traffic channel for the stream 364 and when no burst is received 366, it assumes that the user terminal has closed the stream and closes the stream at its end 367. The base station does not send an express acknowledgement and the user terminal does not wait for an acknowledgment before it closes the stream. Accordingly, the user terminal is ready to open a new stream or handover to another base station as soon as it sends its last data burst. The base station is ready to reallocate the channel to a different stream as soon as it receives no data burst from the user terminal. The channel is then available immediately to be assigned to a different, or the same user for a new stream. As a result of not sending an acknowledgment, in addition to freeing the channel resource sooner, the potential cross-channel interference that might be caused by an acknowledgment signal is also avoided.

[0066] In a lossy channel, the base station can wait through one more frame, listening for the user terminal's data burst. If two frames go by with no data burst from

the user terminal, then the terminal can close the stream. This accounts for the possibility that the user terminal's signal was lost in the channel and was not received by the base station even though the user terminal made a transmission. In such a case, the base station can respond to the reception of no data burst with a data invalid message together with the shut-down request. When the base station again receives no data burst from the user terminal, it will close the stream.

[0067] The possibility of failing to receive a transmitted burst can depend on more than just losses in the channel. It can depend on the received SINR, the transmit power of the other terminal and the type of modulation that has been applied to the signal. Signal modulation such as BPSK (binary phase shift keying) typically have less average power than more complicated modulations such as 16QAM (Quadrature Amplitude Modulation) and so is harder to distinguish from the background noise. The base station can measure, for example, SINR, received signal strength and other characteristics of the received signal. The base station also knows the type of signal modulation that is being used. In one embodiment, several different signal modulation classes are used in the system depending on the channel characteristics and these are communicated in the RA/AA process using a parameter called modclass. Using this information, the base station can determine the possibility that it was not able to differentiate a transmitted signal from the background noise and interference. Using this possibility, the base station can select how many frames it will wait through without being able to detect a signal before it closes the stream.

[0068] In one embodiment, there are several classes of modulation and lower classes of modulation, e.g. BPSK are assigned to noisier channels while higher classes of modulation, e.g. 32QAM are assigned to cleaner channels. In addition the lower

modulation classes transmit with less average power than the higher modulation classes. As a result, there is a correlation between the likelihood of distinguishing the channel from background noise and the modulation class that has been assigned to the stream. The base station, and the user terminal, can use this correlation to select threshold numbers of frames to listen on before closing the stream. For the highest modulation classes, a terminal will listen to the next burst and if no signal is received, it will close the stream. For the lowest modulation classes, the terminal can wait three or four frames before closing the stream.

[0069] Stream Closing Decision

[0070] Figure 2 shows a flow chart for making a decision to close a stream. Each of the features of the decision tree can be used independently of each other or all together, as illustrated here. In addition, as mentioned above, the base station and the user terminal can both close a stream, so either terminal can implement the decision process discussed below. Typically every terminal includes a stream resource manager module that checks buffers, looks for flags and maintains registers as is required to implement the functions described. As mentioned above, a decision to close a stream is normally started when the transmit buffer begins to be emptied 402. The stream resource manager can monitor the transmit buffer and use the bit transfer rate to predict in advance when the buffer will be emptied. This allows sufficient time to handle the other steps shown in Figure 2.

[0071] If the transmit buffer is empty or about to be emptied, the stream resource manager can refer to boundary value flags that are included in the data or received

separately from higher layers 404. The boundary values help to delimit packets and can be used not only for stream closing decisions but also on the receive end to reassemble received packets. Typically, packet size will not match data burst size, especially in a system with variable modulation schemes. In that case, the bit rate of the data stream can vary during the stream or at least in each new stream. Accordingly, the boundary values can be used to reassemble data packets that are transmitted during several different bursts. The boundary values can be in the form of a header for each packet that includes delimiting values among other things or as flag bytes between packets. The boundary values can be restructured for transmission by a higher layer.

[0072] Alternately, the stream resource manager, rather than analyzing the data in the transmit buffer, can receive packet delimiting information from higher layers. By determining packet limits the stream resource manager can analyze the transmit buffer content and determine whether the entire packet has been received in the buffer. If the entire packet has been received in the buffer, then the stream can be closed. If a portion of the packet is yet to be received, then the stream should be kept open to send that portion of the packet that is yet to come. In addition to simply delimiting packet boundaries, the data or the higher layers can also identify the type of data traffic. Streaming video, for example, will require that a data stream be kept open until the video stream is stopped or has ended. Voice telephony may allow the stream to be closed during silent periods on one terminal or the other. Internet browsing may have rather long periods during which no data is transmitted so streams can be closed more aggressively.

[0073] The stream resource manager will typically have a latency timer which must elapse before a stream is closed. The latency timer can be set by a system or

resource administrator or assigned to a terminal by a base station. The latency timer prevents a stream from being closed too quickly. If the timer has not elapsed, i.e. if the transmit buffer has not been empty long enough, then the stream will not be closed 406. A minimum value for the timer is the time required to close a stream and then open a new stream. If the stream is closed before this much time elapses, then, if new data arrives immediately after the stream is closed, a new stream must immediately be opened and the stream opening and closing process does not result in increasing the availability of channels for other users. In addition, it increases latency for the user that closed the stream because it cannot send the data until the new stream is opened. The latency timer might also include an additional increment of time in order to reduce traffic on the random access channel that is used to open streams. It may also include an additional increment of time to account for known latency in other parts of the system. For example, if the data feeding into the buffer was subject to random delays, the length of these delays can be added so that such a delay is not interpreted as the end of the data stream. Further aspects of the latency timer are discussed below.

[0074] After the latency timer has timed out, the stream resource manager refers to a predictive timer 408. This timer may be referred to before or as an alternative to the latency timer. If the predictive timer has not timed out, then the stream will not be closed. The stream resource manager establishes the predictive timer by looking at the history of receiving data in the transmit buffer. This history includes identifying packet delimiters. The history will show what the typical gaps are between packets and what the typical delays are between packets. The delays may be a part of the communications channels being used but it may also be caused by the habits of the terminal user. The

stream resource manager can look at all of these trends and establish a predictive timer value that optimizes efficiency based on the patterns of the particular user.

[0075] If the predictive timer has also timed out, then the terminal sends a shut-down request 410 as discussed above. It then listens to the return link. If it receives an implicit acknowledgement of its shut-down request 412, i.e. if the other terminal stops transmitting, then it closes the stream 414.

[0076] Operation

[0077] As can be understood from the discussion above, the present invention allows a circuit-switched communications channel to be opened with a single handshake by using a previous virtual circuit to set up constant routing information. The channel is the stream opened by sending a RA-rts message or an in-band rts message. The virtual circuit is the result of registration. All the routing information that does not depend on the physical channel assigned for the stream is done one time up front for all channels. As a result, the handshake for opening a new stream that will allocate a physical channel can be minimized to a single channel request (RA-request to send) and a single channel grant (AA-clear to send). As can be seen from the discussion above, both of these messages are short, reducing their impact on the random access channel. In addition, the two-message protocol minimizes the time needed to open physical channels.

[0078] As described above, the present invention also minimizes the overhead in closing a stream. This allows the channel to be reused by other users as quickly as possible. It also allows the stream closing strategy to be more aggressive in choosing when to close streams. As described above, in one embodiment, when one side of the air

interface is ready to close a stream, it will send a message to its peer requesting that that the stream be shut down. If the peer is also ready to close the stream, then the peer will shut down its side of the stream immediately. The original side will sense that the peer is no longer using the channel, and will take this as an implicit acknowledgement of the shut down request. Thus, there is a single-pass handshake, and furthermore, the response part of the handshake is implicit. Making the response implicit allows the peer side to avoid using one last frame to send an explicit acknowledgement message.

[0079] The fast stream request and grant exchange allows the system to aggressively close channels when there is no data to send or receive. This allows a channel that would otherwise be allocated but idle to be available as soon as possible to another user that will actually use it. When one side of an air-interface channel has no data to send and also receives no data for a configurable number of frames, it will attempt to deallocate the channel. This configurable number of frames constitutes a timer and can be set to a small value, including a single frame, so the channel will be deallocated as soon as there is no data to send and receive.

[0080] Accordingly, a part of the aggressive channel deallocation scheme is to have the ability to hold channels open even when there is no data to send or receive. This allows the stream to kept open long enough to bridge the gap between the time when there is no old data to send and receive on the stream, and the time when new data arrives from higher layers. This is done using, for example, the packet boundary delimiters and a stream manager that communicates with higher layers. Without the stream manager, the aggressive stream deallocation strategy would deallocate the first stream as soon as no data were available, and then have to allocate a second stream for the new data. Holding

the stream open saves the overhead of deallocating the first stream and allocating the second at the cost of leaving the stream idle until the new data arrives.

[0081] As an alternative to the stream manager, the timer mentioned above can be set to a slightly larger value based on network traffic patterns, so the stream will not be deallocated quite as quickly giving new data has a chance to arrive. This kind of a timer is a type of predictive timer based on history. If one side of an air interface channel is not receiving data, but has only partially received a data unit of a higher protocol layer, it will be more conservative in closing its stream, since it assumes the remainder of the higher layer's data unit will be delivered soon. Also, if the stream is being used for certain diagnostics in which no data is currently being transferred but there is a certainty of later data transfer, the stream closing strategy is more conservative.

[0082] The configurable number of frames can be chosen with consideration of the break-even point between a purely aggressive deallocation strategy and a conservative deallocation strategy that waits for new data to arrive. The break-even point will depend on the network traffic patterns and may vary from user to user. Typically, it is the point at which the sum of channel deallocation and allocation latencies is the same as the time the channel would be idle if it were not deallocated. Therefore, a reasonable compromise is to hold a channel open even when there is no data for the same amount of time that it would take the user to deallocate the old channel and then allocate a new channel.

[0083] The latency of allocating a new channel can be seen as having two parts. First, there is the time from requesting a stream to the assignment of the stream, the RA-rts/AA-cts process discussed above. Second, there is the time from assignment to when the channel reaches full bandwidth utilization. In the present invention, this second

latency portion is kept very short because spatial parameters and channel characteristics are already well-developed during the registration process and are improved through the channel request/assignment process. If the break-even point is used to set the timer then, in the case that new data arrived at the latest time at which the channel is still kept open, the system loading efficiency is the same as if the channel had been immediately deallocated and then reallocated. However, the latency seen by the user is still decreased by the channel allocation time. In one embodiment of the present invention, this timer is set by default to the approximate sum of stream closing and stream opening times. This is a latency timer as discussed above.

[0084] As mentioned above, the selection of this timer can be improved using expected data traffic patterns. This can be done by using separate latency and predictive timers as discussed above with respect to Figure 2 or the timers can be combined. To the extent to which future data traffic patterns can be predicted using, for example, past data patterns or external information, the deallocation strategy can be further optimized to conserve system resources using the variable predictive timer. Streams can be closed quickly if it is predicted that no new data will arrive soon. User latency can be improved by closing streams conservatively if it is predicted that new data will soon arrive.

[0085] Base Station Structure

[0086] The present invention relates to wireless communication systems and may be a fixed-access or mobile-access wireless network using spatial division multiple access (SDMA) technology in combination with multiple access systems, such as time division multiple access (TDMA), frequency division multiple access (FDMA) and code

division multiple access (CDMA). Multiple access can be combined with frequency division duplexing (FDD) or time division duplexing (TDD). Figure 3 shows an example of a base station of a wireless communications system or network suitable for implementing the present invention. The system or network includes a number of subscriber stations, also referred to as remote terminals or user terminals, such as that shown in Figure 4. The base station may be connected to a wide area network (WAN) through its host DSP 231 for providing any required data services and connections external to the immediate wireless system. To support spatial diversity, a plurality of antennas 103 is used, for example four antennas, although other numbers of antennas may be selected.

[0087] The outputs of the antennas are connected to a duplexer switch 107, which in this TDD system is a time switch. Two possible implementations of switch 107 are as a frequency duplexer in a frequency division duplex (FDD) system, and as a time switch in a time division duplex (TDD) system. When receiving, the antenna outputs are connected via switch 107 to a receiver 205, and are mixed down in analog by RF receiver ("RX") modules 205 from the carrier frequency to an FM intermediate frequency ("IF"). This signal then is digitized (sampled) by analog to digital converters ("ADCs") 209. Final down-converting to baseband is carried out digitally. The down-converting can be done using finite impulse response (FIR) filtering techniques. This is shown as block 213. The invention can be adapted to suit a wide variety of RF and IF carrier frequencies and bands.

[0088] There are, in the present example, four down-converted outputs from each antenna's digital filter device 213, one per receive timeslot. The particular number of timeslots can be varied to suit network needs. While the present example uses four

uplink and four downlink timeslots for each TDD frame, desirable results have also been achieved with three timeslots for the uplink and downlink in each frame. For each of the four receive timeslots, the four down-converted outputs from the four antennas are fed to a digital signal processor (DSP) device 217 (hereinafter "timeslot processor") for further processing, including calibration, according to one aspect of this invention. Four Motorola DSP56311 DSPs can be used as timeslot processors, one per receive timeslot. The timeslot processors 217 monitor the received signal power and estimate the frequency offset and time alignment. They also determine smart antenna weights for each antenna element. These are used in the spatial division multiple access scheme to determine a signal from a particular remote user and to demodulate the determined signal.

[0089] The output of the timeslot processors 217 is demodulated burst data for each of the four receive timeslots. This data is sent to the host DSP processor 231 whose main function is to control all elements of the system and interface with the higher level processing, which is the processing which deals with what signals are required for communications in all the different control and service communication channels defined in the system's communication protocol. The host DSP 231 can be a Motorola DSP56303. In addition, timeslot processors send the determined receive weights for each user terminal to the host DSP 231. The host DSP 231 maintains state and timing information, receives uplink burst data from the timeslot processors 217, and programs the timeslot processors 217. In addition it decrypts, descrambles, checks error detecting code, and deconstructs bursts of the uplink signals, then formats the uplink signals to be sent for higher level processing in other parts of the base station. With respect to the other parts of the base station it formats service data and traffic data for further higher processing in the base station, receives downlink messages and traffic data from the other

parts of the base station, processes the downlink bursts and formats and sends the downlink bursts to a transmit controller/modulator, shown as 237. The host DSP also manages programming of other components of the base station including the transmit controller/modulator 237 and the RF timing controller shown as 233.

[0090] The RF timing controller 233 interfaces with the RF system, shown as block 245 and also produces a number of timing signals that are used by both the RF system and the modem. The RF controller 233 reads and transmits power monitoring and control values, controls the duplexer 107 and receives timing parameters and other settings for each burst from the host DSP 231.

[0091] The transmit controller/modulator 237, receives transmit data from the host DSP 231, four symbols at a time. The transmit controller uses this data to produce analog IF outputs which are sent to the RF transmitter (TX) modules 245. Specifically, the received data bits are converted into a complex modulated signal, up-converted to an IF frequency, 4-times over-sampled, multiplied by transmit weights obtained from host DSP 231, and converted via digital to analog converters ("DACs") which are part of transmit controller/modulator 237 to analog transmit waveforms. The analog waveforms are sent to the transmit modules 245. The transmit modules 245 up-convert the signals to the transmission frequency and amplify the signals. The amplified transmission signal outputs are sent to antennas 103 via the duplexer/time switch 107.

[0092] User Terminal Structure

[0093] Figure 4 depicts an example component arrangement in a user terminal that provides data or voice communication. The user terminal's antenna 45 is connected

to a duplexer 46 to permit antenna 45 to be used for both transmission and reception. The antenna can be omni-directional or directional. For optimal performance, the antenna can be made up of multiple elements and employ spatial processing as discussed above for the base station. In an alternate embodiment, separate receive and transmit antennas are used eliminating the need for the duplexer 46. In another alternate embodiment, where time division duplexing is used, a transmit/receive (TR) switch can be used instead of a duplexer as is well-known in the art. The duplexer output 47 serves as input to a receiver 48. The receiver 48 produces a down-converted signal 49 which is the input to a demodulator 51. A demodulated received sound or voice signal 67 is input to a speaker 66.

[0094] The user terminal has a corresponding transmit chain in which data or voice to be transmitted is modulated in a modulator 57. The modulated signal to be transmitted 59, output by the modulator 57, is up-converted and amplified by a transmitter 60, producing a transmitter output signal 61. The transmitter output 61 is then input to the duplexer 46 for transmission by the antenna 45.

[0095] The demodulated received data 52 is supplied to a user terminal central processing unit 68 (CPU) as is received data before demodulation 50. The user terminal CPU 68 can be implemented with a standard DSP (digital signal processor) device such as a Motorola series 56300 DSP. This DSP can also perform the functions of the demodulator 51 and the modulator 57. The user terminal CPU 68 controls the receiver through line 63, the transmitter through line 62, the demodulator through line 52 and the modulator through line 58. It also communicates with a keyboard 53 through line 54 and a display 56 through line 55. A microphone 64 and speaker 66 are connected through the modulator 57 and the demodulator 51 through lines 65 and 66, respectively for a voice

communications user terminal. In another embodiment, the microphone and speaker are also in direct communication with the CPU to provide voice or data communications.

[0096] The user terminal's voice signal to be transmitted 65 from the microphone 64 is input to a modulator 57. Traffic and control data to be transmitted 58 is supplied by the user terminal's CPU 68. Control data 58 is transmitted to base stations during registration, session initiation and termination as well as during the session or stream.

[0097] In an alternate embodiment, the speaker 66, and the microphone 64 are replaced or augmented by digital interfaces well-known in the art that allow data to be transmitted to and from an external data processing device (for example, a computer). In one embodiment, the user terminal's CPU is coupled to a standard digital interface such as a PCMCIA interface to an external computer and the display, keyboard, microphone and speaker are a part of the external computer. The user terminal's CPU 68 communicates with these components through the digital interface and the external computer's controller. For data only communications, the microphone and speaker can be deleted. For voice only communications, the keyboard and display can be deleted.

[0098] In the description above, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

[0099] The present invention includes various steps. The steps of the present invention may be performed by hardware components, such as those shown in Figure 3 or 4, or may be embodied in machine-executable instructions, which may be used to

cause a general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

[00100] The present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of media or machine-readable medium suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

[00101] Importantly, while the present invention has been described in the context of a wireless internet data system for portable handsets, it can be applied to a wide variety of different wireless systems in which data is exchanged. Such systems include voice, video, music, broadcast and other types of data systems without external connections. The present invention can be applied to fixed user terminals as well as to low and high mobility terminals. Many of the methods are described herein in a basic form but steps can be added to or deleted from any of the methods and information can be added or subtracted from any of the described messages without departing from the basic scope of the present invention. It will be apparent to those skilled in the art that many further modifications and adaptations can be made. The particular embodiments are not

provided to limit the invention but to illustrate it. The scope of the present invention is not to be determined by the specific examples provided above but only by the claims below.